



# Beyond CS106A

Chris Piech  
CS106A, Stanford University

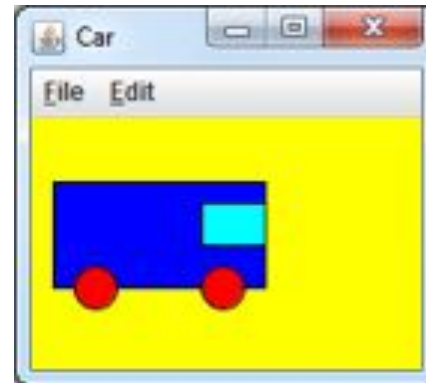
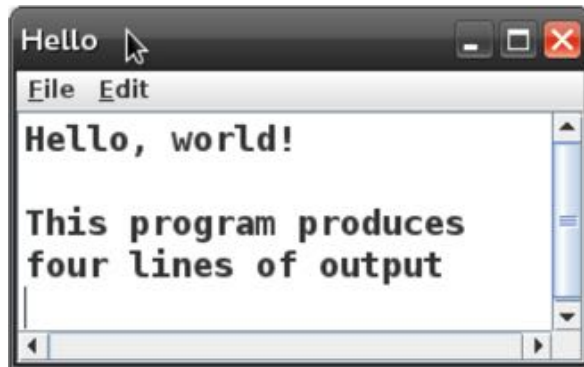





Life after ACM

# Life After The ACM Libraries

- All quarter we have relied on the **ACM Java libraries**.
  - Karel, ConsoleProgram, RandomGenerator
  - GraphicsProgram, GOval, GRect, GOval, GLine, GImage, ...



- Today we will see how **standard Java** programs are made.



# Using the ACM Libraries

```
import acm.program.*;

public class MyProgram extends ConsoleProgram {
    public void run() {
        println("Hello, world!");
    }
}
```

- This is a console program written using the ACM libraries.
  - It uses the **ConsoleProgram** class to represent a console.
  - The **run** method contains the program code.
  - The **println** method prints output to the graphical console.



# A Barebones Java Program

```
public class Hello {  
    public static void main(String[] args) {  
        System.out.println("Hello, world!");  
    }  
}
```

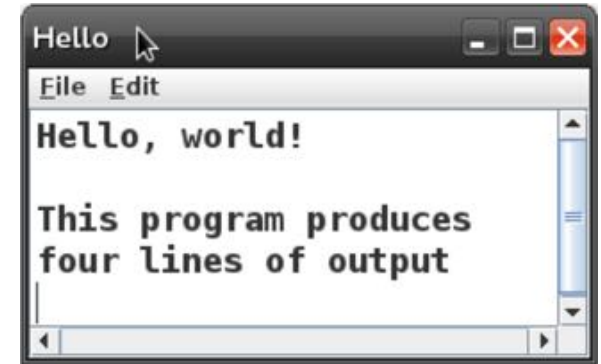
- The method **main** is the true entry point for a Java program.
  - It must have the exact heading shown above.
  - The **String[] args** are "command line arguments" (ignored).
  - The **println** command's true name is **System.out.println**.
  - Standard Java methods are **static** unless part of a class of objects.



Lets make Hello World!

# Console Programs

- What does the **ConsoleProgram** library class do?
  - Creates a new graphical **window**
  - Puts a scrollable **text area** into it
  - Provides `print` and `println` commands to send text **output** to that window
  - contains a **main method** that calls your program class's `run` method
    - `ConsoleProgram`'s `run` is empty, but you extend and override it



```
public class Hello extends ConsoleProgram {  
    public void run() {  
        println("Hello, world!");  
    }  
}
```





# ACM console input

```
public class Age extends ConsoleProgram {
    public void run() {
        String name = readLine("What's your name? ");
        int age = readInt("How old are you? ");
        int years = 65 - age;
        println(name + " has " + years
            + " years until retirement!");
    }
}
```

- The ACM library has simple console input commands like `readLine`, `readInt`, `readDouble`, and so on.
- These methods display a 'prompt' message, wait for input, re-prompt if the user types a bad value, and return the input.



# Java console input

```
public class Age {
    public static void main(String[] args) {
        Scanner console = new Scanner(System.in);
        System.out.print("What's your name? ");
        String name = console.nextLine();
        System.out.print("How old are you? ");
        int age = console.nextInt();
        int years = 65 - age;
        System.out.println(name + " has " + years
            + " years until retirement!");
    }
}
```

- In standard Java, you must create a Scanner or similar object to read input from the console, which is also called `System.in`.
  - It does not automatically re-prompt and can crash on bad input.



# Graphics Programs

The ACM library does several things to make graphics easier:

- Automatically creates and displays a **window** on the screen.
  - In standard Java, we must do this ourselves; it is called a JFrame.
- Sets up a **drawing canvas** in the center of the window  
In standard Java, we must create our own drawing canvas.
- Provides convenient methods to listen for mouse events.
  - In standard Java, event handling takes a bit more code to set up.



# ACM GUI example

```
public class ColorFun extends Program {
    public void init() {
        JButton button1 = new JButton("Red!");
        JButton button2 = new JButton("Blue!");
        add(button1, SOUTH);
        add(button2, SOUTH);
        addActionListeners();
    }

    public void actionPerformed(ActionEvent event) {
        if (event.getActionCommand().equals("Red!")) {
            setBackground(Color.BLUE);
        } else {
            setBackground(Color.RED);
        }
    }
}
```



# Java GUI example

```
public class ColorFun implements ActionListener {
    public static void main(String[] args) {
        new ColorFun().init();
    }
    private JFrame frame;
    public void init() {
        frame = new JFrame("ColorFun");
        frame.setSize(500, 300);
        JButton button1 = new JButton("Red!");
        JButton button2 = new JButton("Blue!");
        button1.addActionListener(this);
        button2.addActionListener(this);
        frame.add(button1, "South");
        frame.add(button2, "South");
        frame.setVisible(true);
    }
    public void actionPerformed(ActionEvent event) {
        if (event.getActionCommand().equals("Red!")) {
            frame.setBackground(Color.BLUE);
        } else {
            frame.setBackground(Color.RED);
        }
    }
}
```



# Summary

- **Benefits of libraries:**

- simplify syntax/rough edges of language/API
- avoid re-writing the same code over and over
- possible to make advanced programs quickly
- leverage work of others



- **Drawbacks of libraries:**

- limitations on usage; e.g. ACM library cannot be re-distributed for commercial purposes



# Programming Languages



# Java

```
ArrayList<Double> evens = new ArrayList<>();  
for(int i = 0; i < 100; i++) {  
    if(i % 2 == 0) {  
        evens.add(i);  
    }  
}  
println(evens);
```

prints [2, 4, 6, 8, 10, 12, ... ]





# C++

```
Vector<double> evens;  
for(int i = 0; i < 100; i++) {  
    if(i % 2 == 0) {  
        evens.add(i);  
    }  
}  
cout << evens << endl;
```

prints [2, 4, 6, 8, 10, 12, ... ]



# Python

```
evens = []  
for i in range(100):  
    if i % 2 == 0:  
        evens.append(i)  
print evens
```

prints [2, 4, 6, 8, 10, 12, ... ]



# Javascript

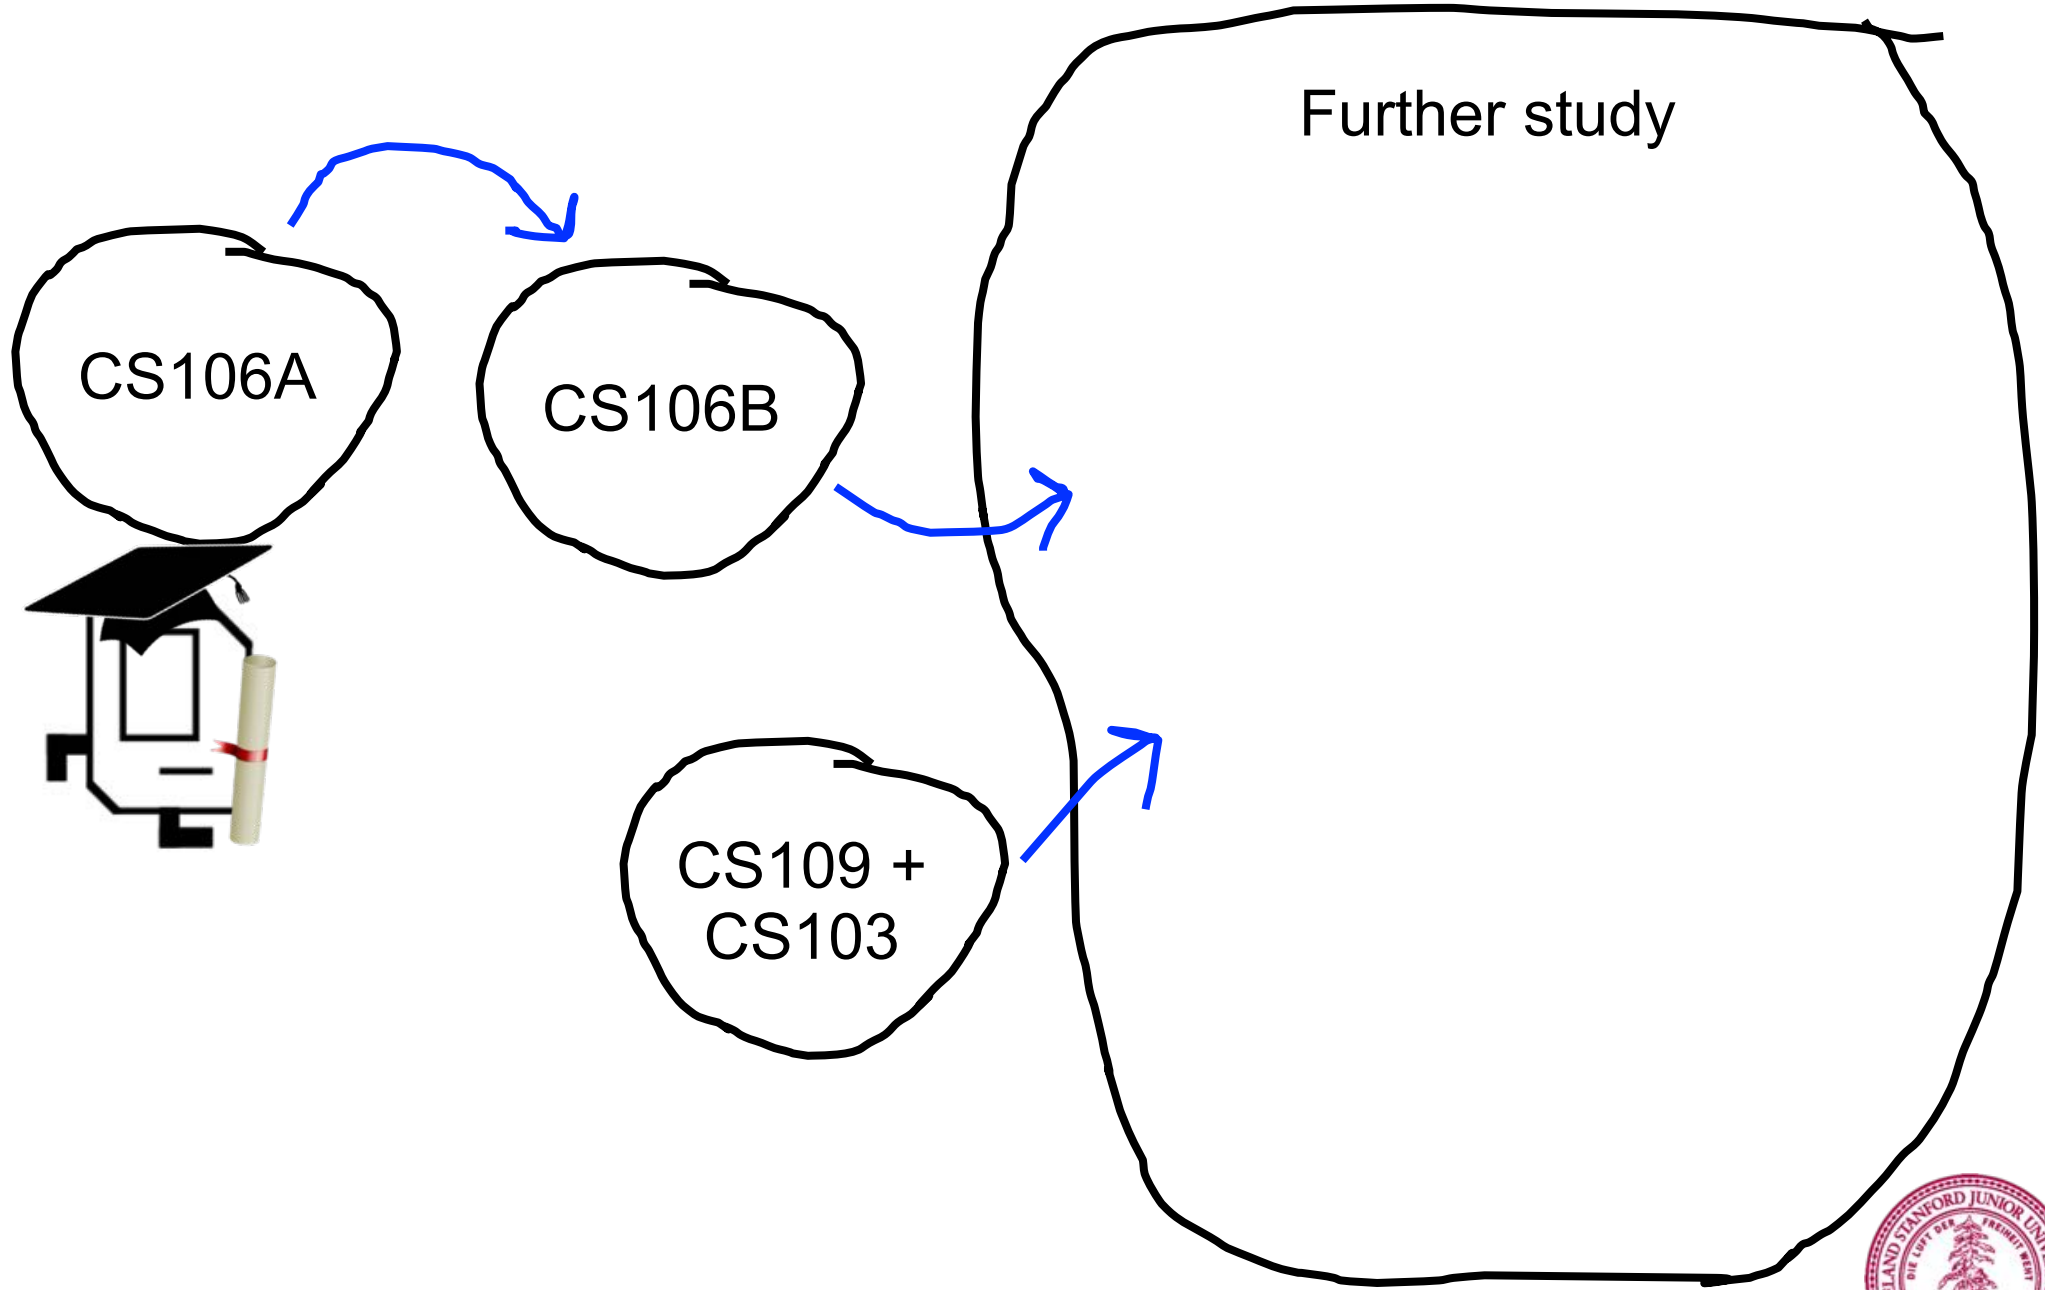
```
var evens = []
for(var i = 0; i < 100; i++) {
    if(i % 2 == 0) {
        evens.push(i)
    }
}
console.log(evens)
```

prints [2, 4, 6, 8, 10, 12, ... ]

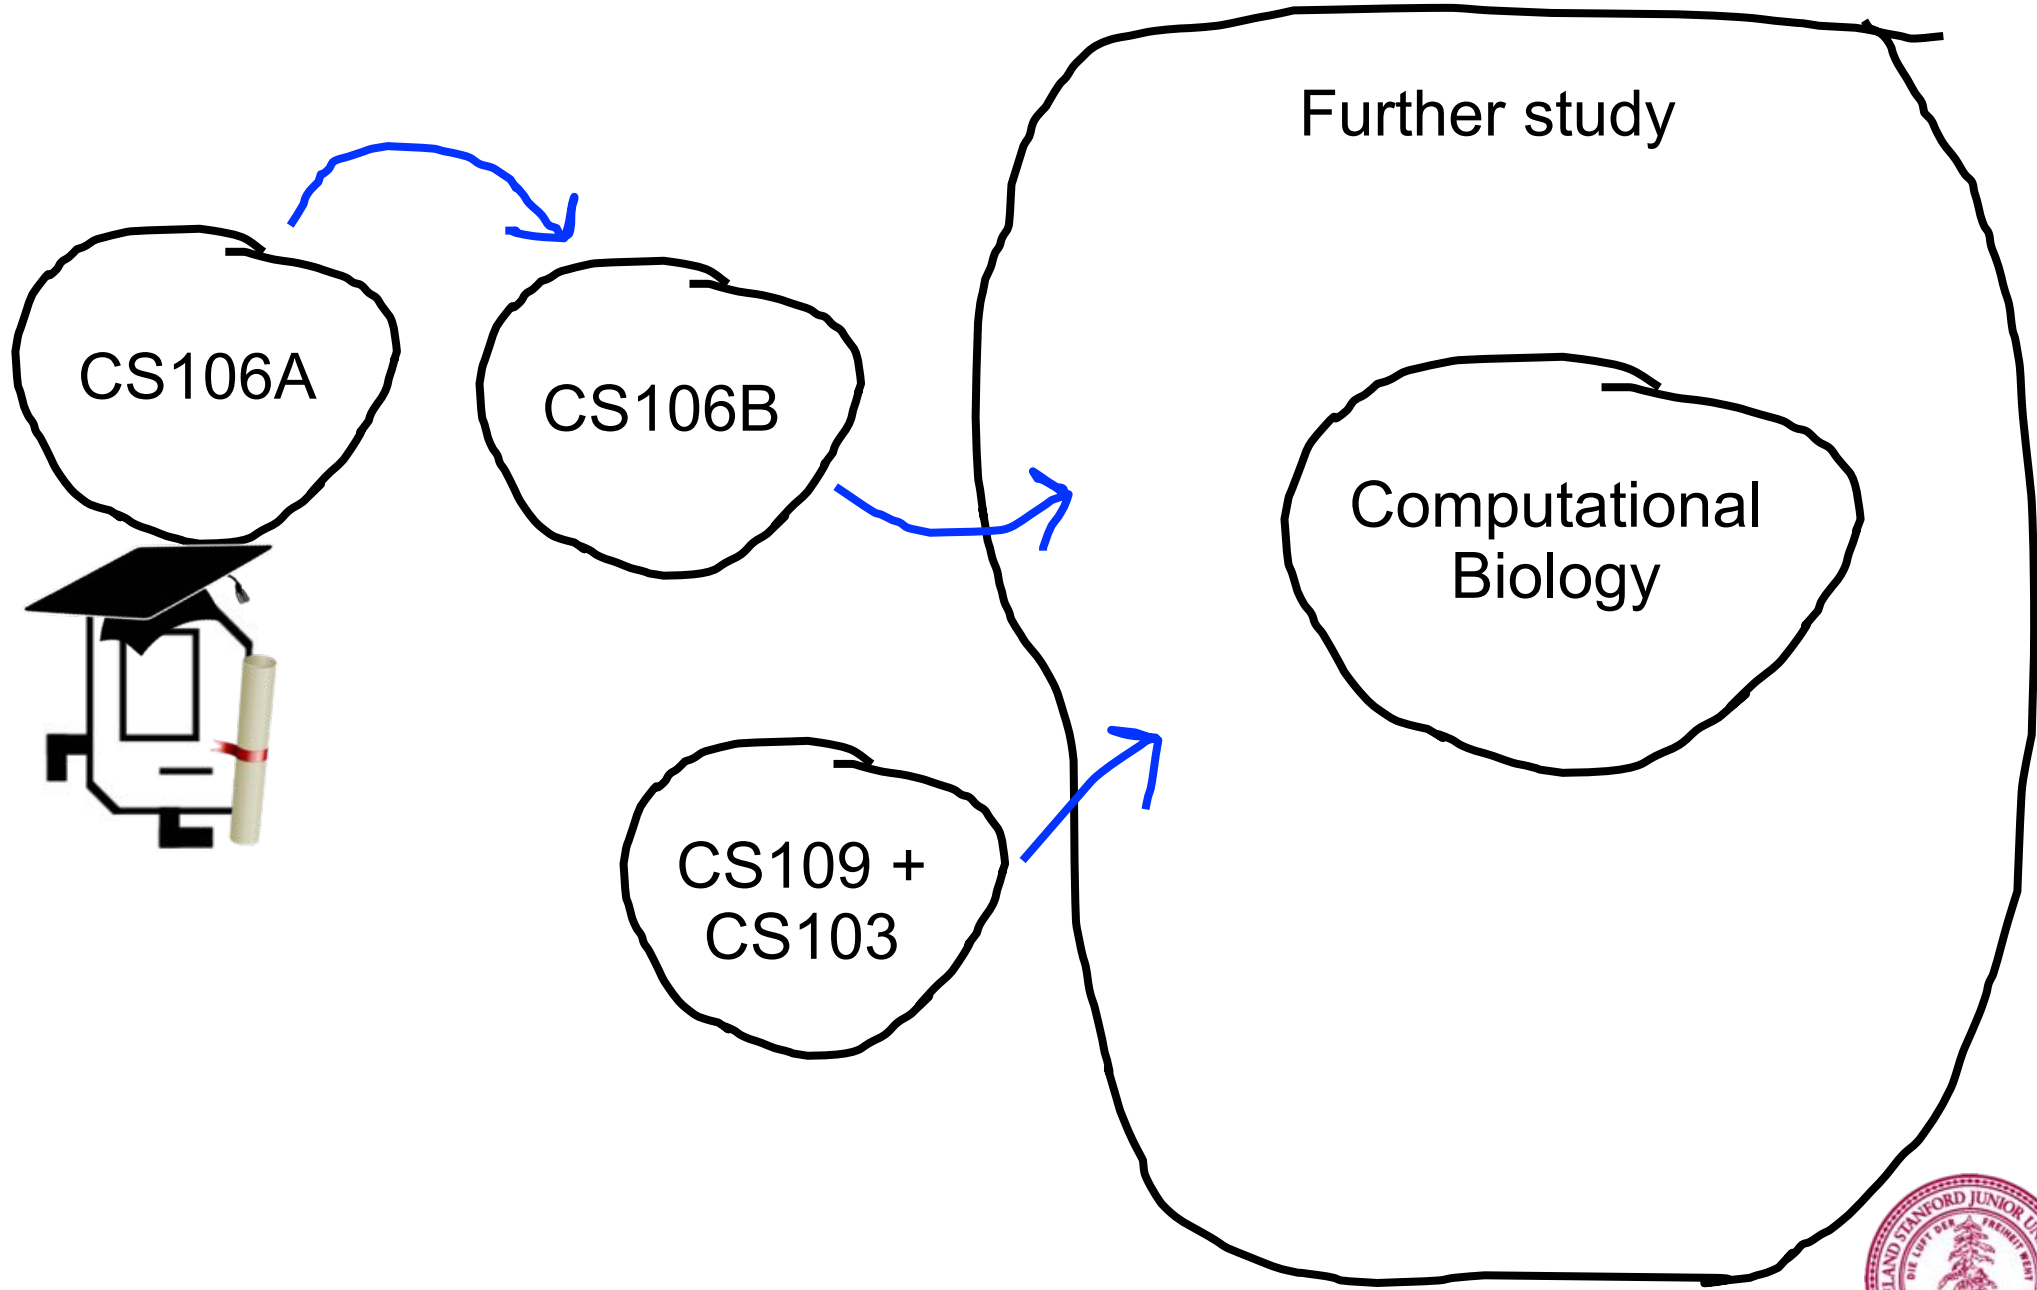


# Future in the CS curriculum

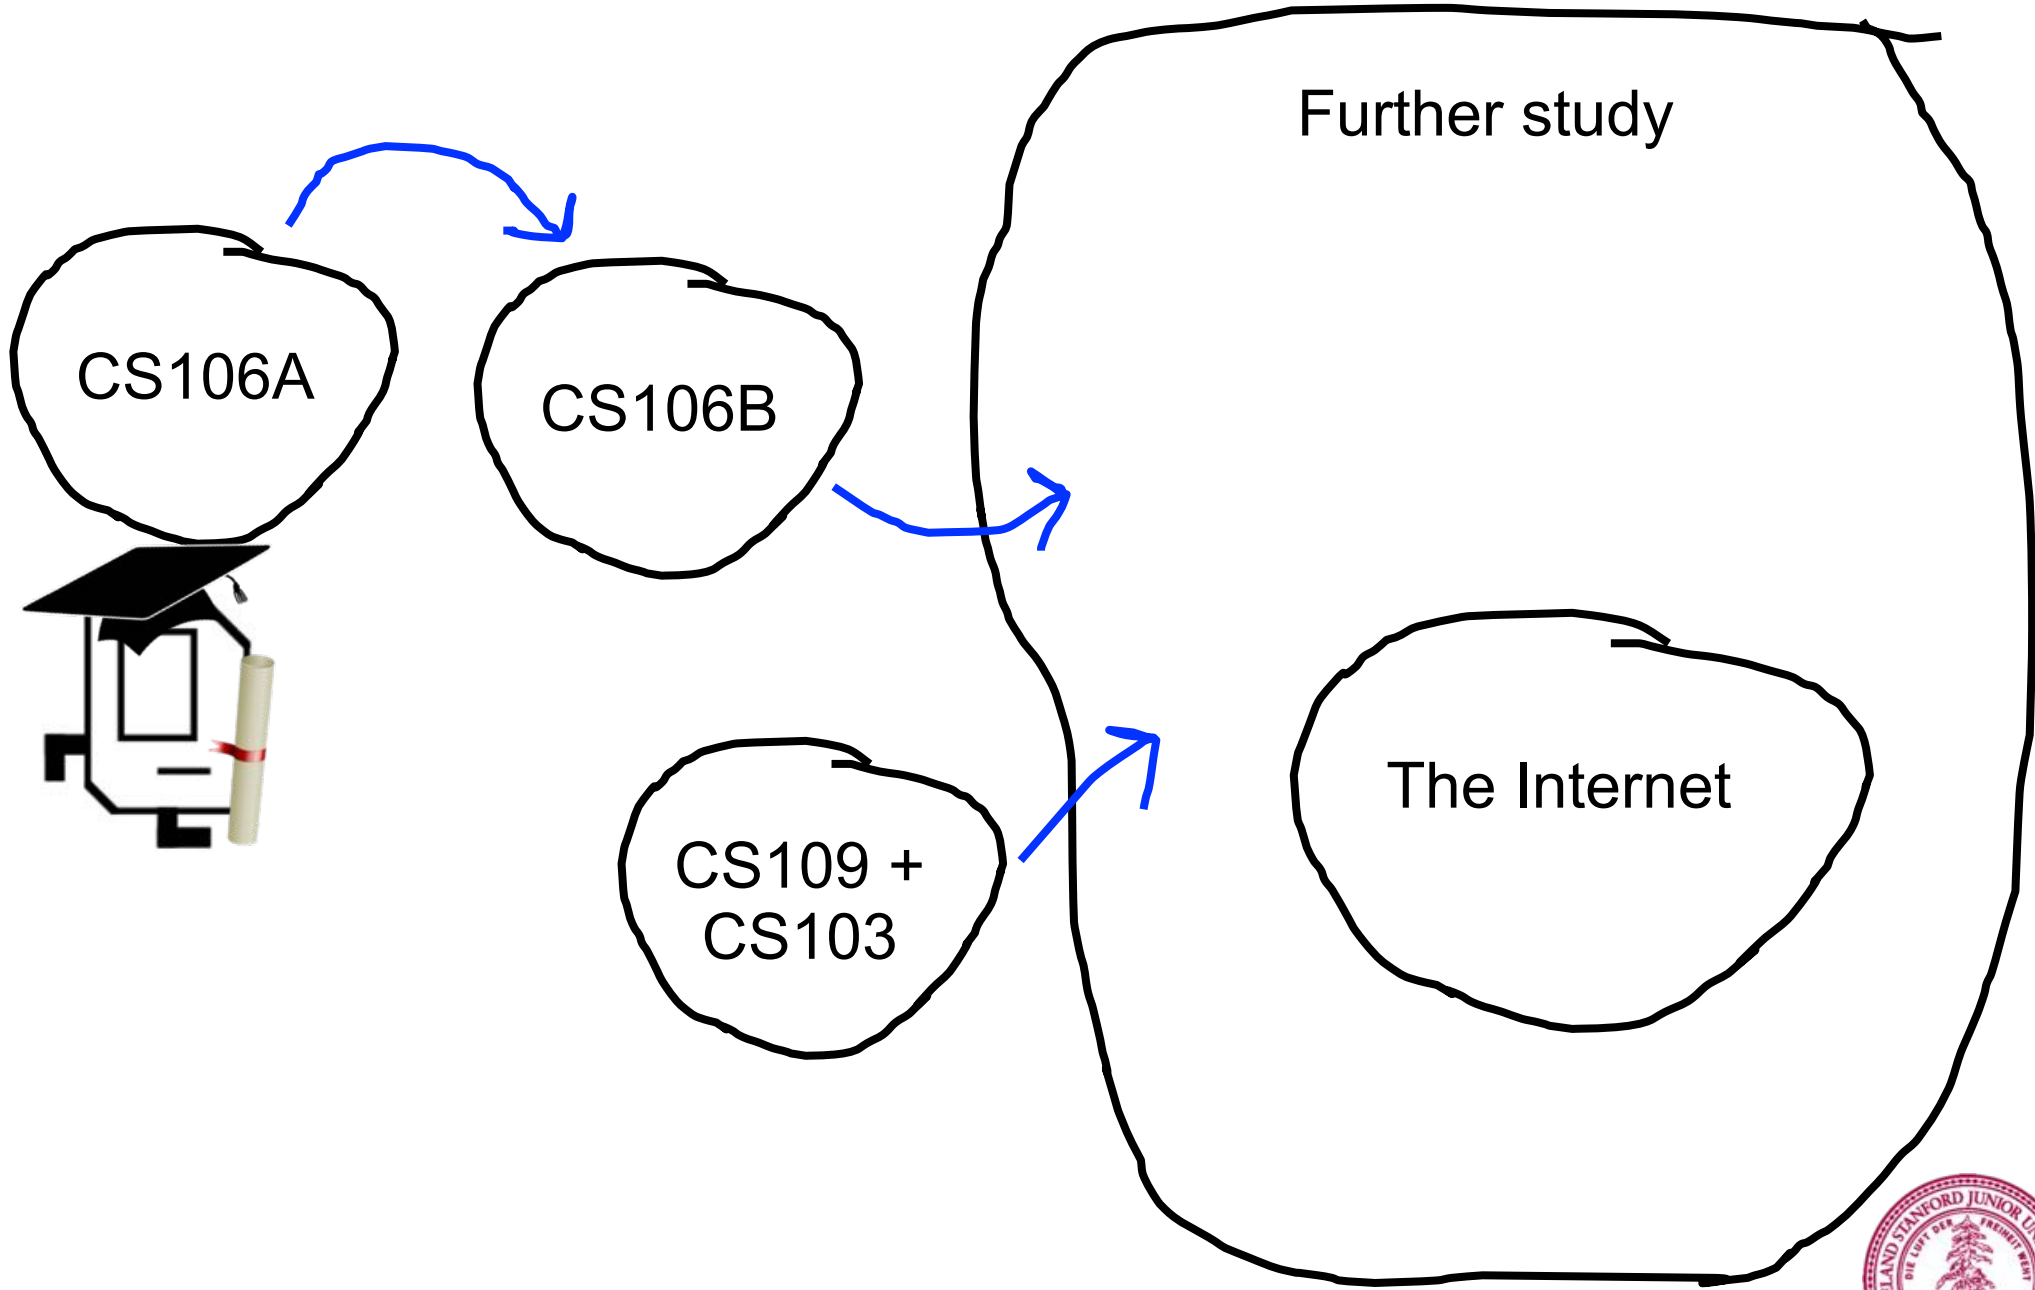
# The CS Curriculum



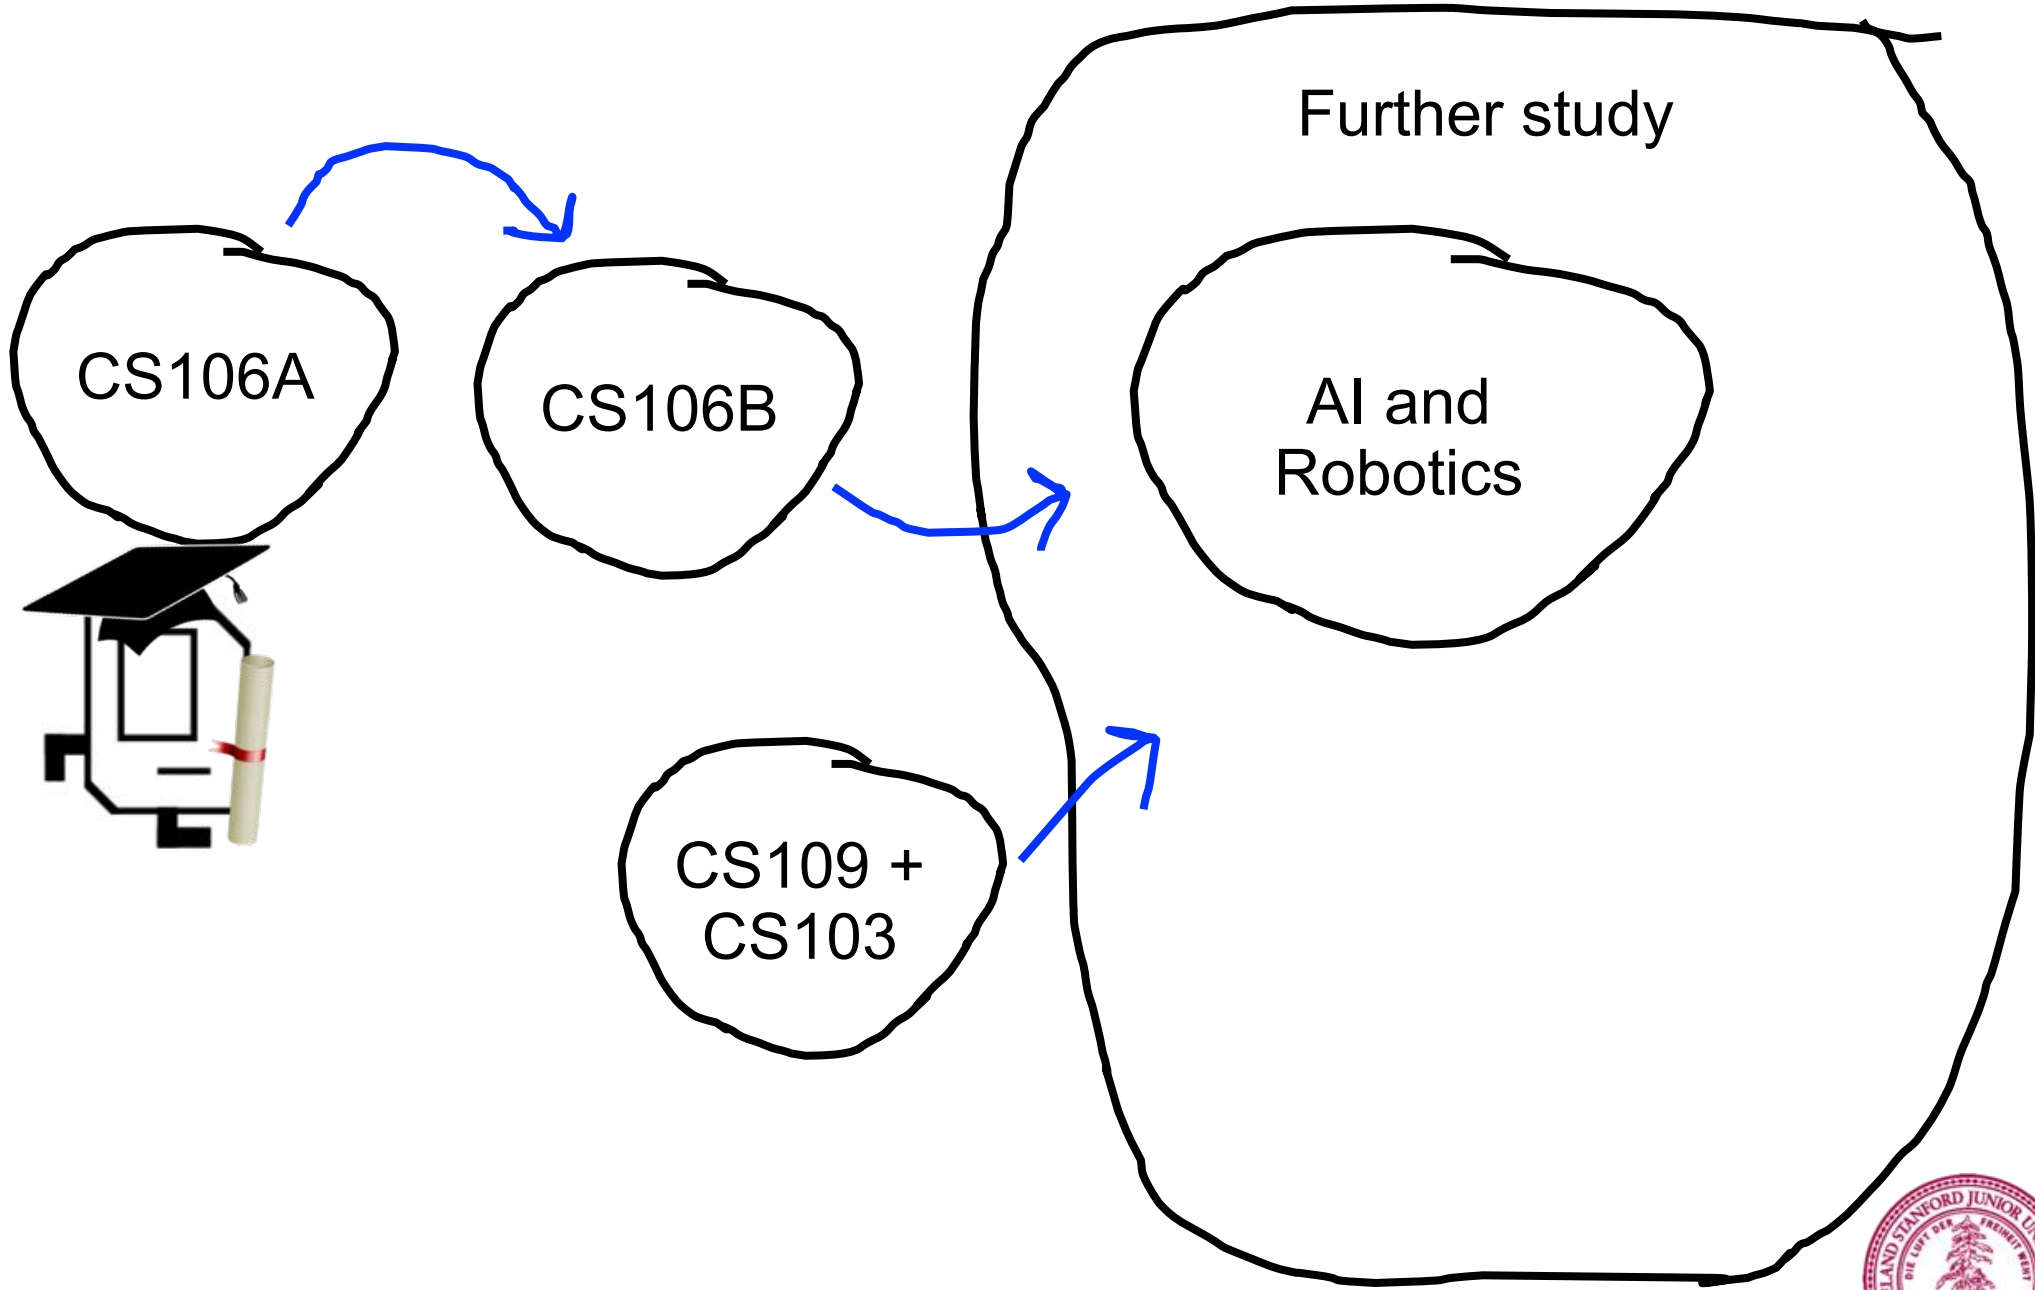
# The CS Curriculum



# The CS Curriculum

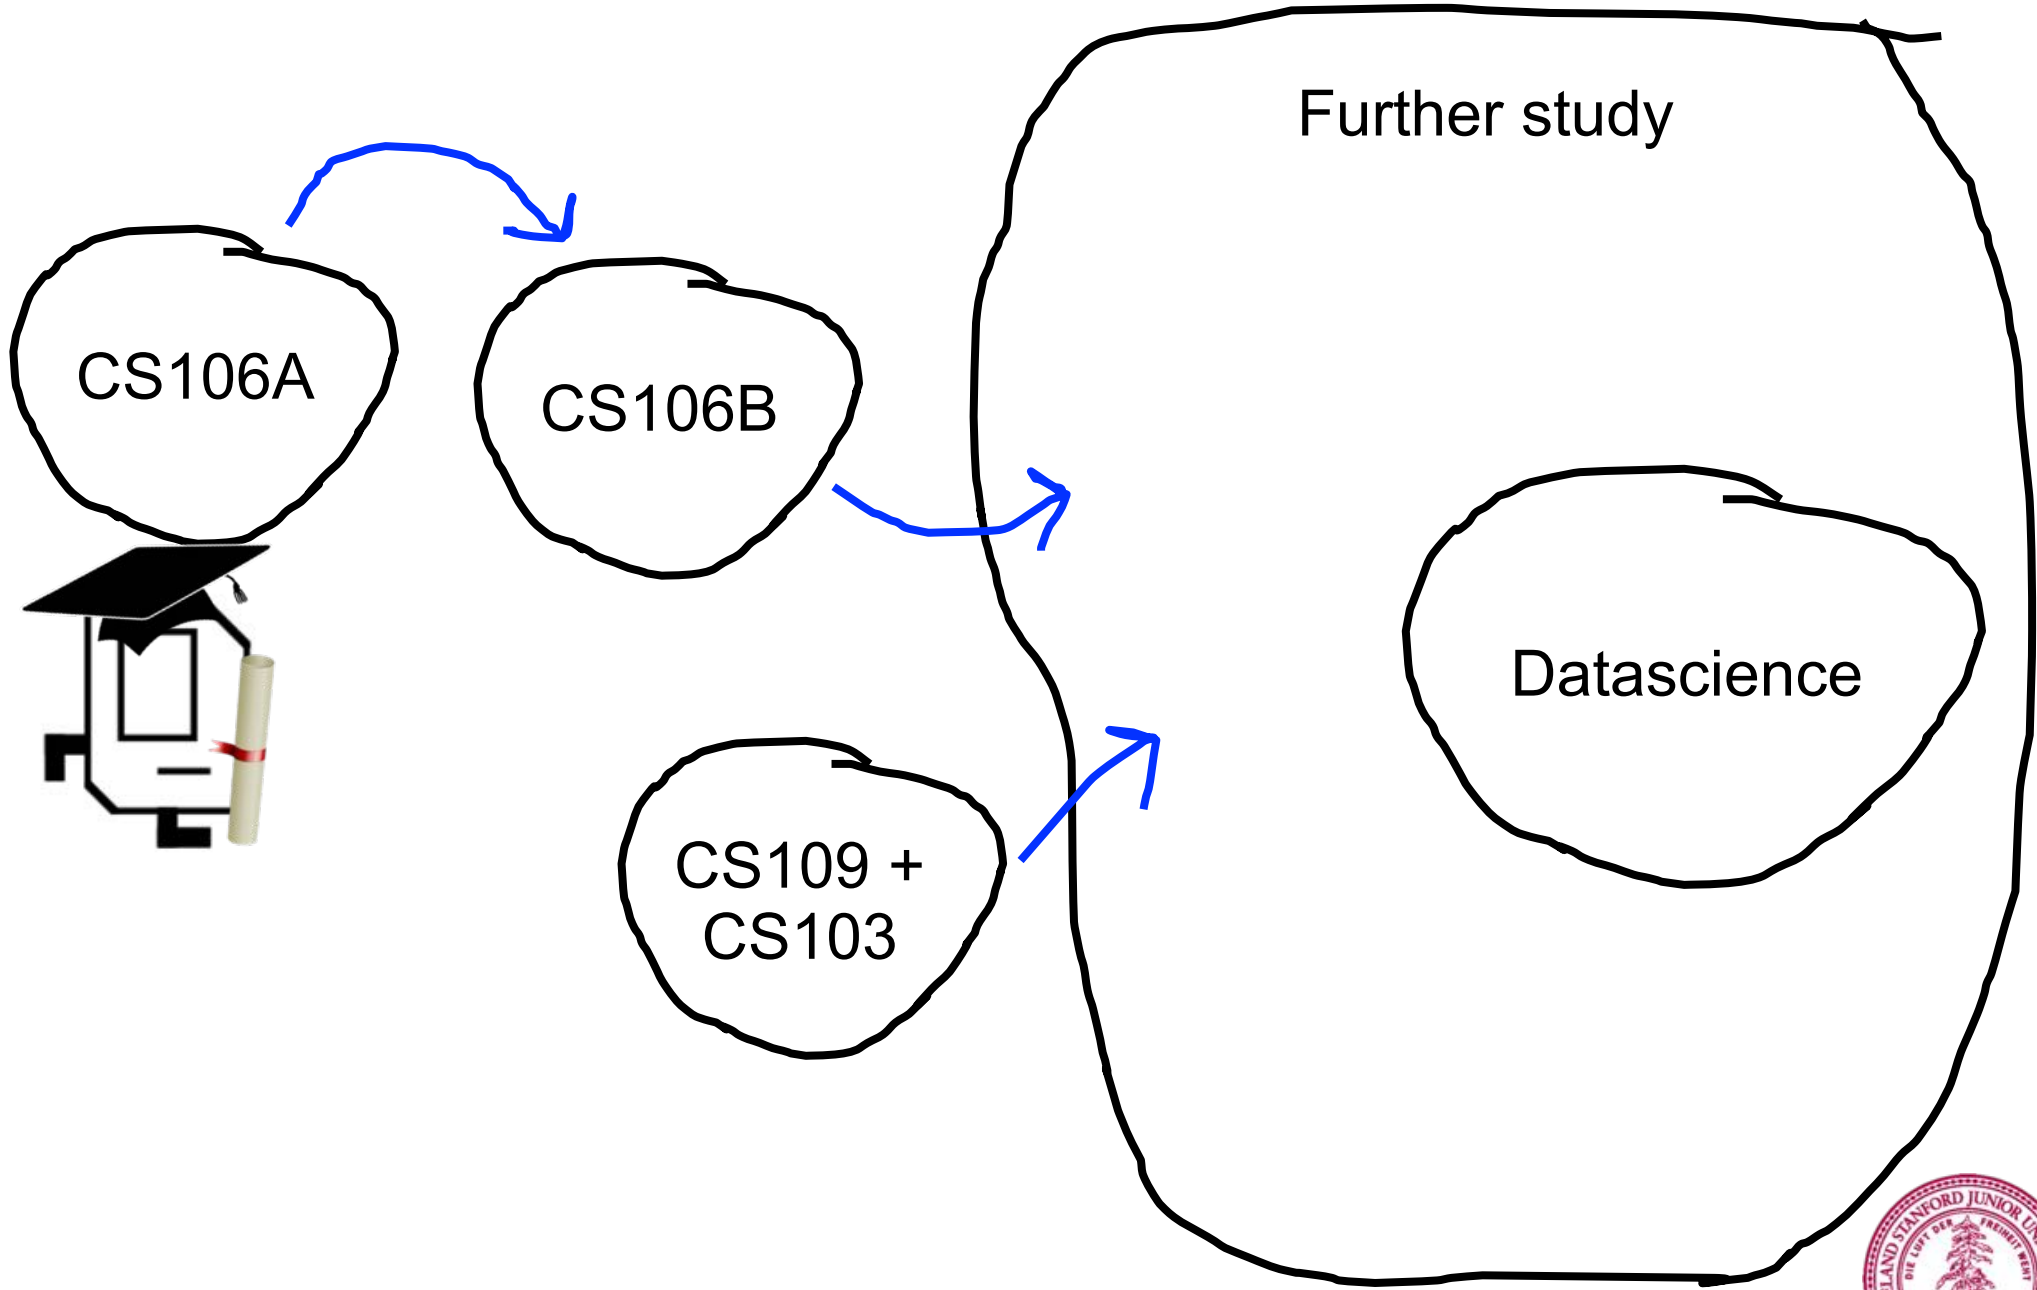


# The CS Curriculum

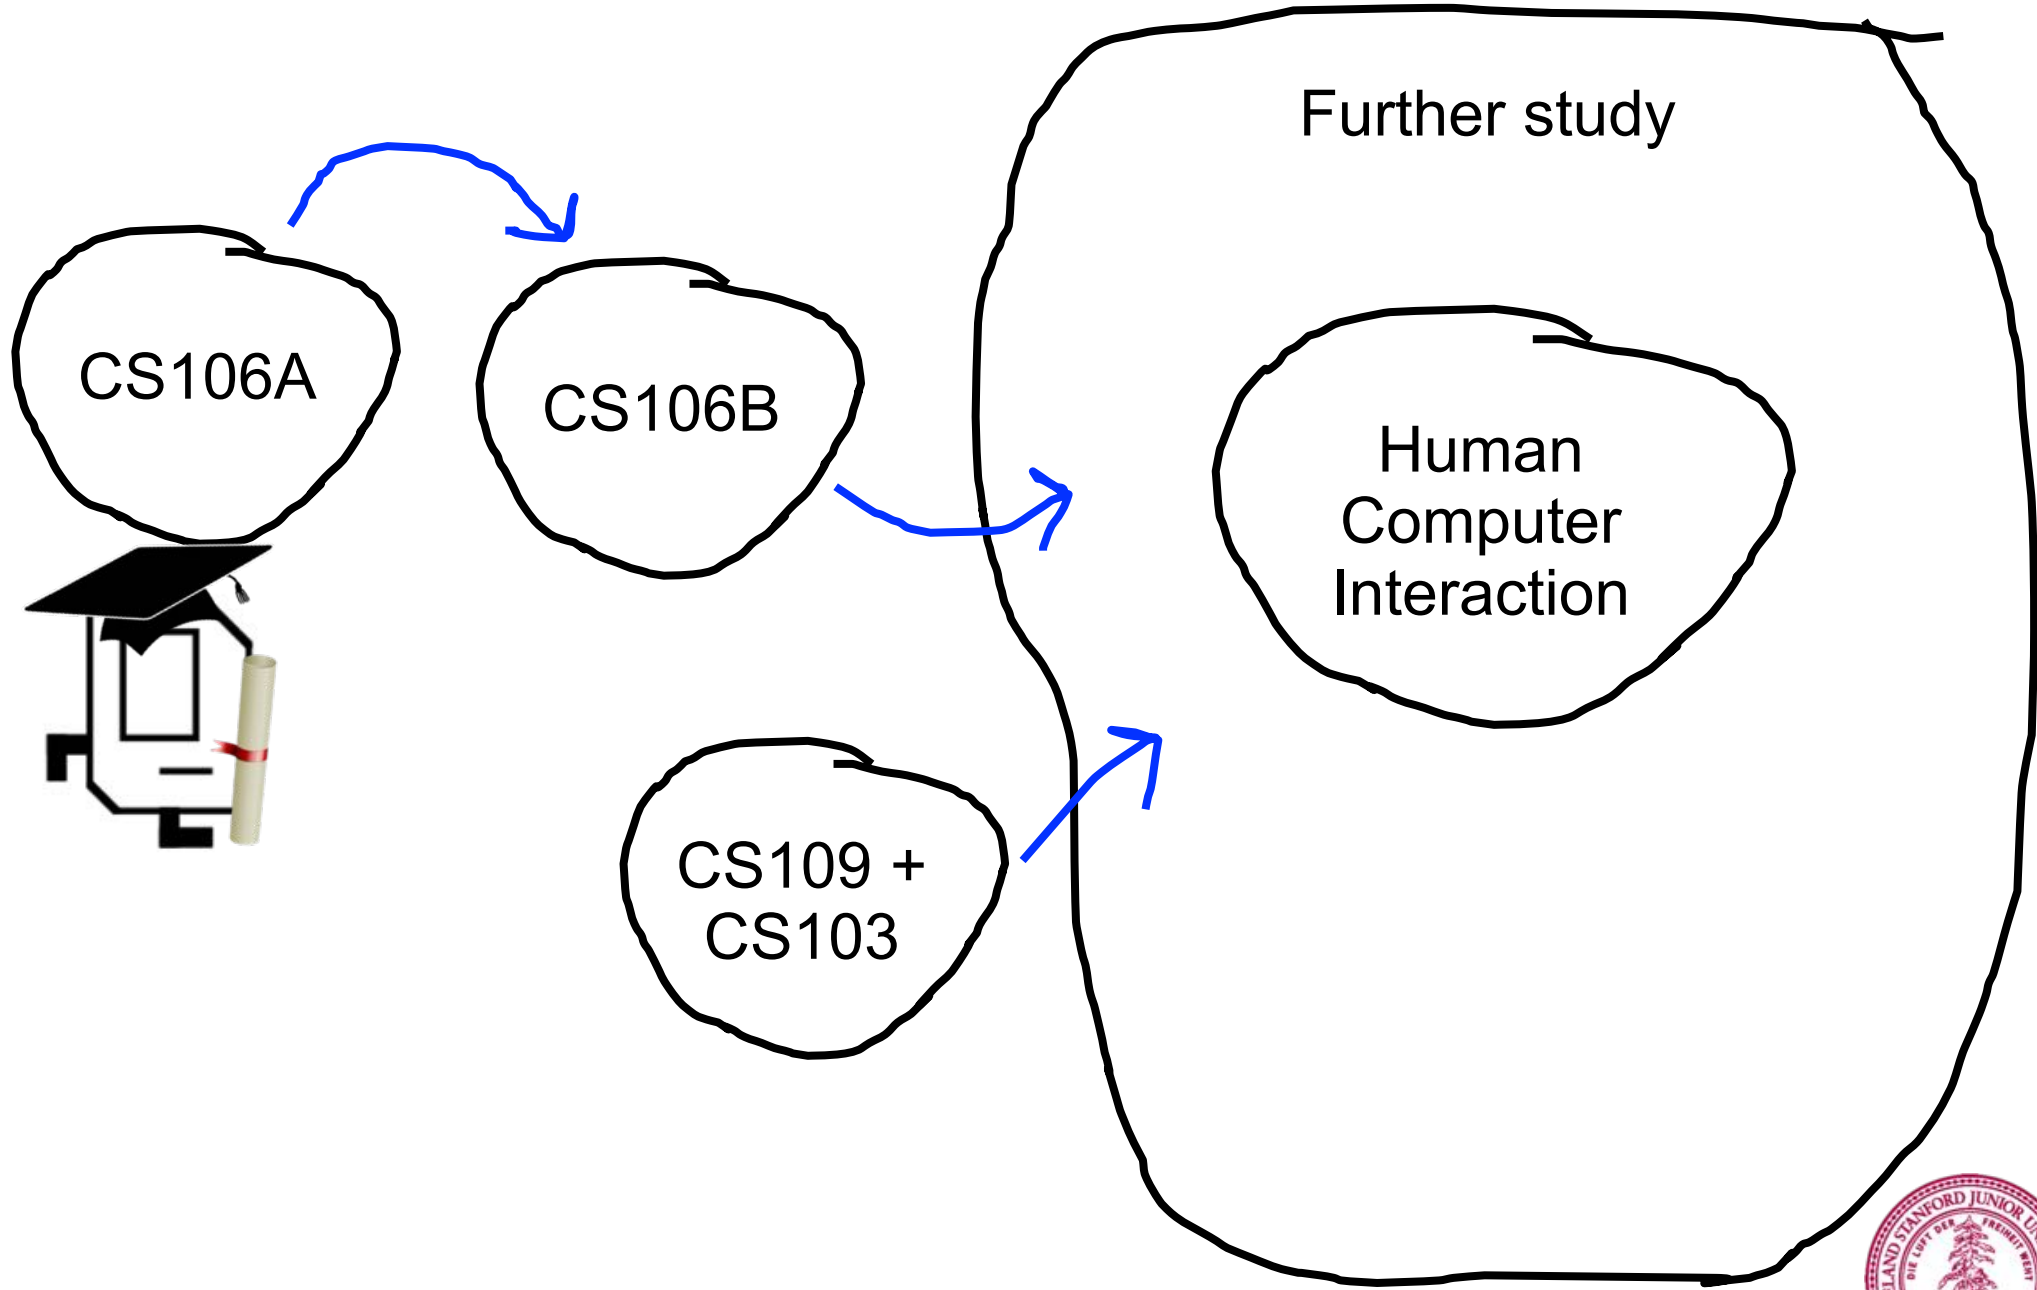




# The CS Curriculum



# The CS Curriculum



# Machine Learning

# Machine Learning or, How we learned to decompose

There is something going on  
in the world of AI

Something big (for us)...

*[suspense]*

# How can we develop intelligent **agents**?





# Volunteer



Computer programs



How can we develop intelligent **agents**?

Better than chance



As well as humans



# Big Milestones Starting in 1997



1997 Deep Blue



2005 Stanley



2011 Watson



# Self Driving Cars





# Computers Making Art



# The Last Remaining Board Game




# Early Optimism 1950

1952



1955



# Early Optimism 1950

“Machines will be capable, within twenty years, of doing any work a man can do.”

–Herbert Simon, 1952





# Underwhelming Results 1950s to 1980s

*The spirit is willing but the flesh is weak.*



(Russian)



*The vodka is good but the meat is rotten.*

The world is too complex



**BRACE YOURSELVES**



**WINTER IS COMING**

# Simple Example: Identifying Cats

- We have a picture and we want to know if it's a cat or not.



→ true



→ false



→ true



→ false





7 Decibel - Tännchen  
Kirkh on deviatora  
21 4 54 11 222  
COMPUTER

# Identifying Cats

Here's one way you might code this...

```
private void isCat(GImage animal) {
    int[][] pixels = animal.getPixelArray();
    if (containsTwoEyes(pixels)){
        if (hasWhiskers(pixels)){
            if (hasPointyEars(pixels)){
                return true;
            }
        }
    }
    return false;
}
```





# Some Tricky Cases



# Pros / Cons

- Pros
  - Matches our human intuition about what a cat is
  - Easy to understand the code
- Cons
  - Requires us to explicitly enumerate every feature that's important, and know how important it is
  - Need to write code to detect eyes, and whiskers, and the pointiness of ears
  - Will never improve... cannot learn from its mistakes



Hard problems seemed impossible.




Great idea #1 learn from experience

# Machine Learning: Learn From Experience




Great idea #2 inspired by biology


# Neuron




# Neuron



# Neuron




# Neuron






# Neuron







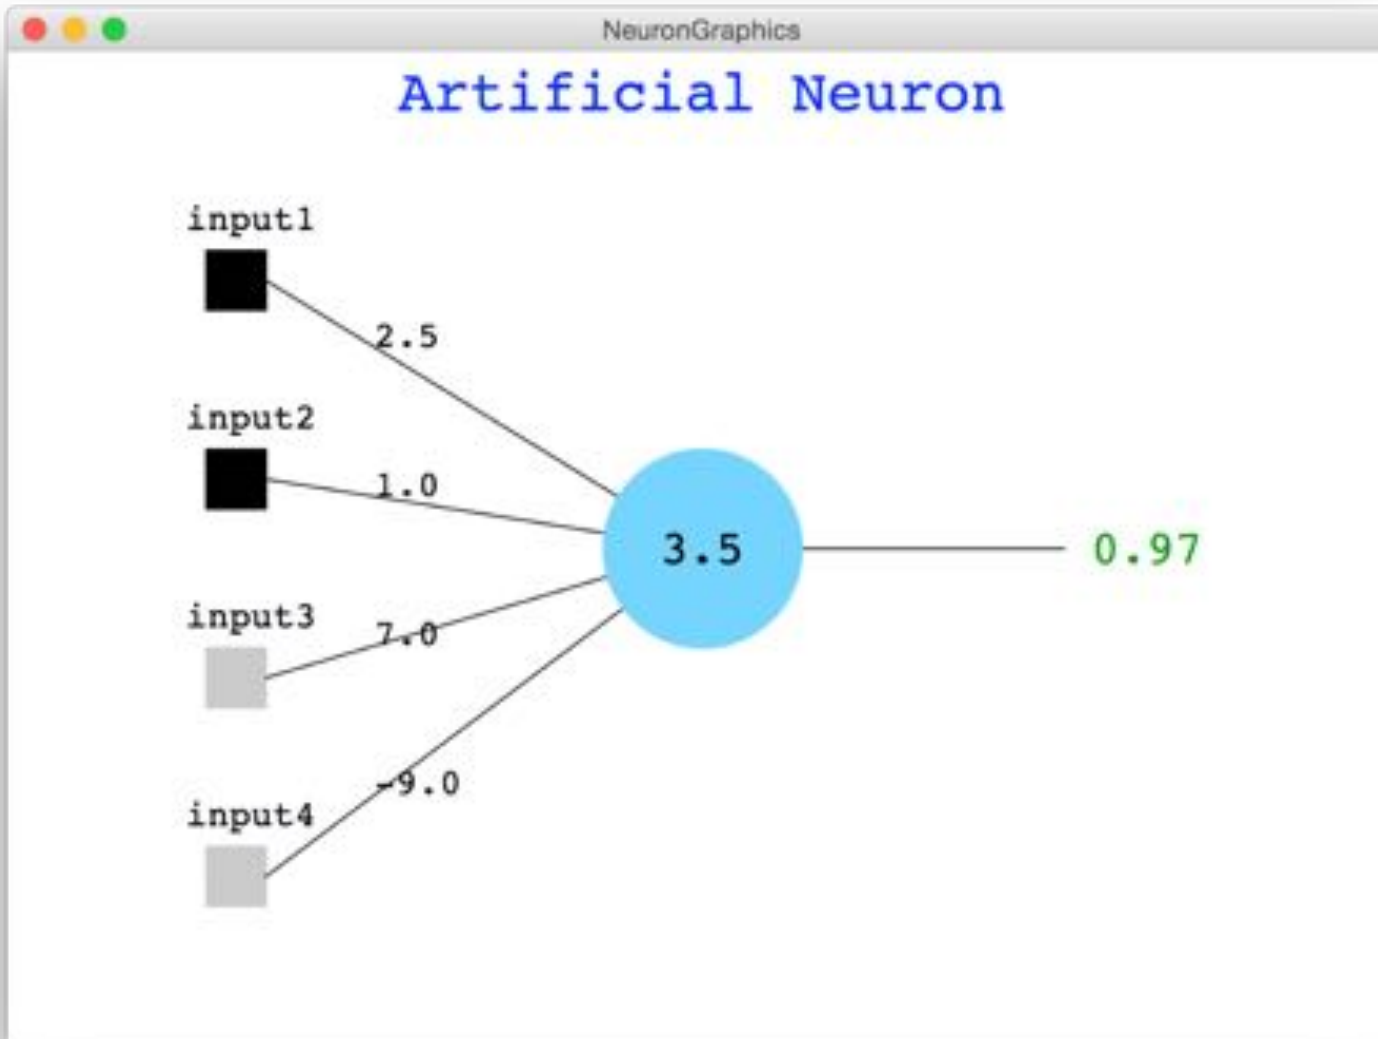
# Some Inputs are More Important



# Artificial Neurons



# Java Demo

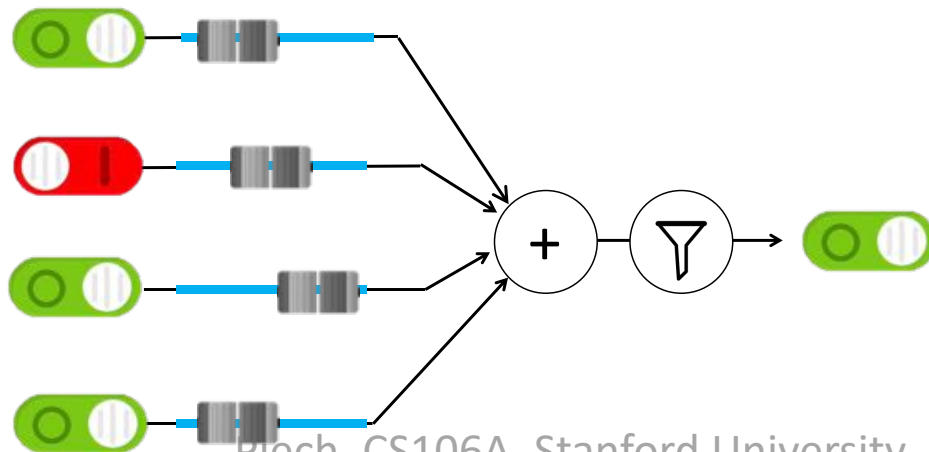


# Artificial Neuron

```
// calculate the activation of a neuron
private double activate(double[] weights, double[] inputs) {
    double weightedSum = 0;
    for(int i = 0; i < inputs.length; i++) {
        weightedSum += weights[i] * inputs[i];
    }


    return squash(weightedSum);
}
```

```
// the sigmoid function forces a value to be between 0 and 1
private double squash(double value) {
    return 1 / (1 + Math.exp(-value));
}
```



# Digit Recognition Example

Let's make feature vectors from pictures of numbers



$$\text{input} = [0, 0, 0, 0, \dots, 1, 0, 0, 1, \dots, 0, 0, 1, 0]$$
$$\text{label} = 0$$



$$\text{input} = [0, 0, 1, 1, \dots, 0, 1, 1, 0, \dots, 0, 1, 0, 0]$$
$$\text{label} = 1$$



# Computer Vision



Piech, CS106A, Stanford University



# Classification



That is a picture  
of a **one**



# Classification



That is a picture  
of a **zero**





# Classification



That is a picture  
of an **zero**





\* It doesn't have to  
be correct all of the  
time



Can you do it?


# Single Neuron




This means it predicts a 0



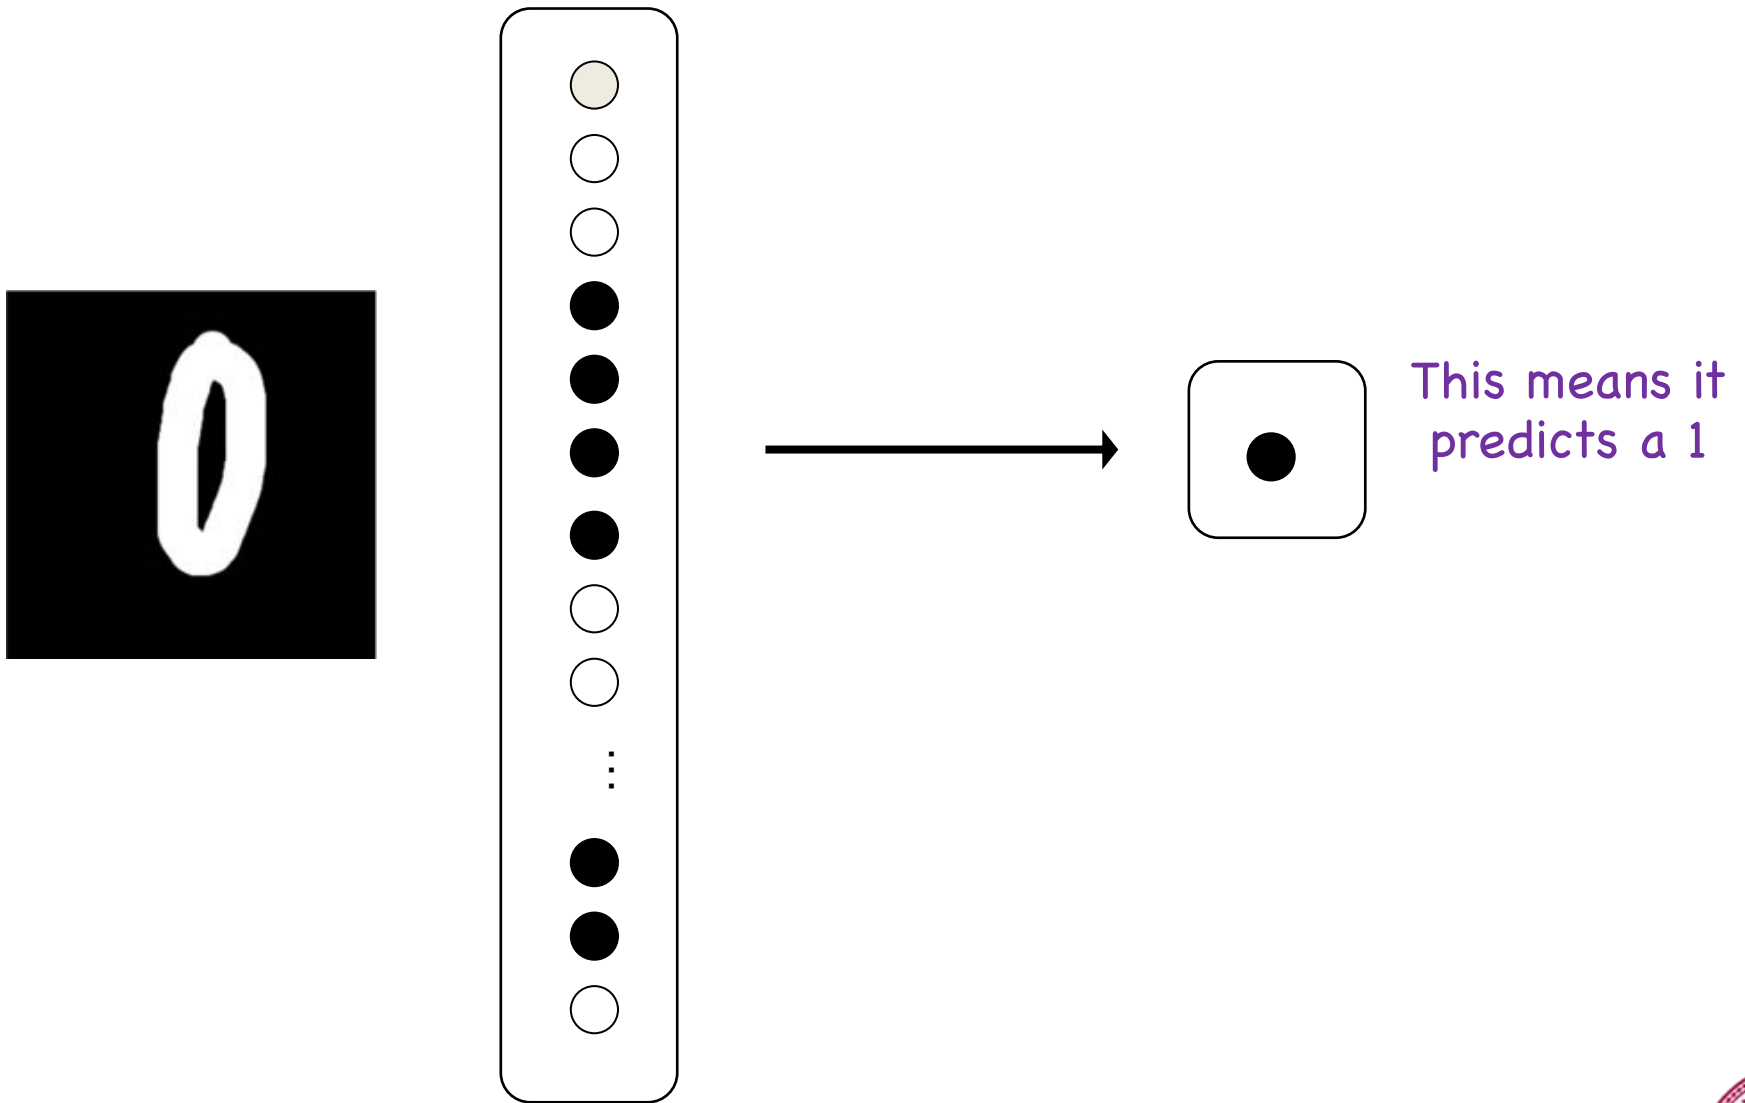
# Single Neuron



# Single Neuron





# Not So Good




# Biological Basis for Neural Networks


- A neuron




- Your brain



Actually, it's probably someone else's brain




# We Can Put Neurons Together








# Demonstration

Draw your number here



X  

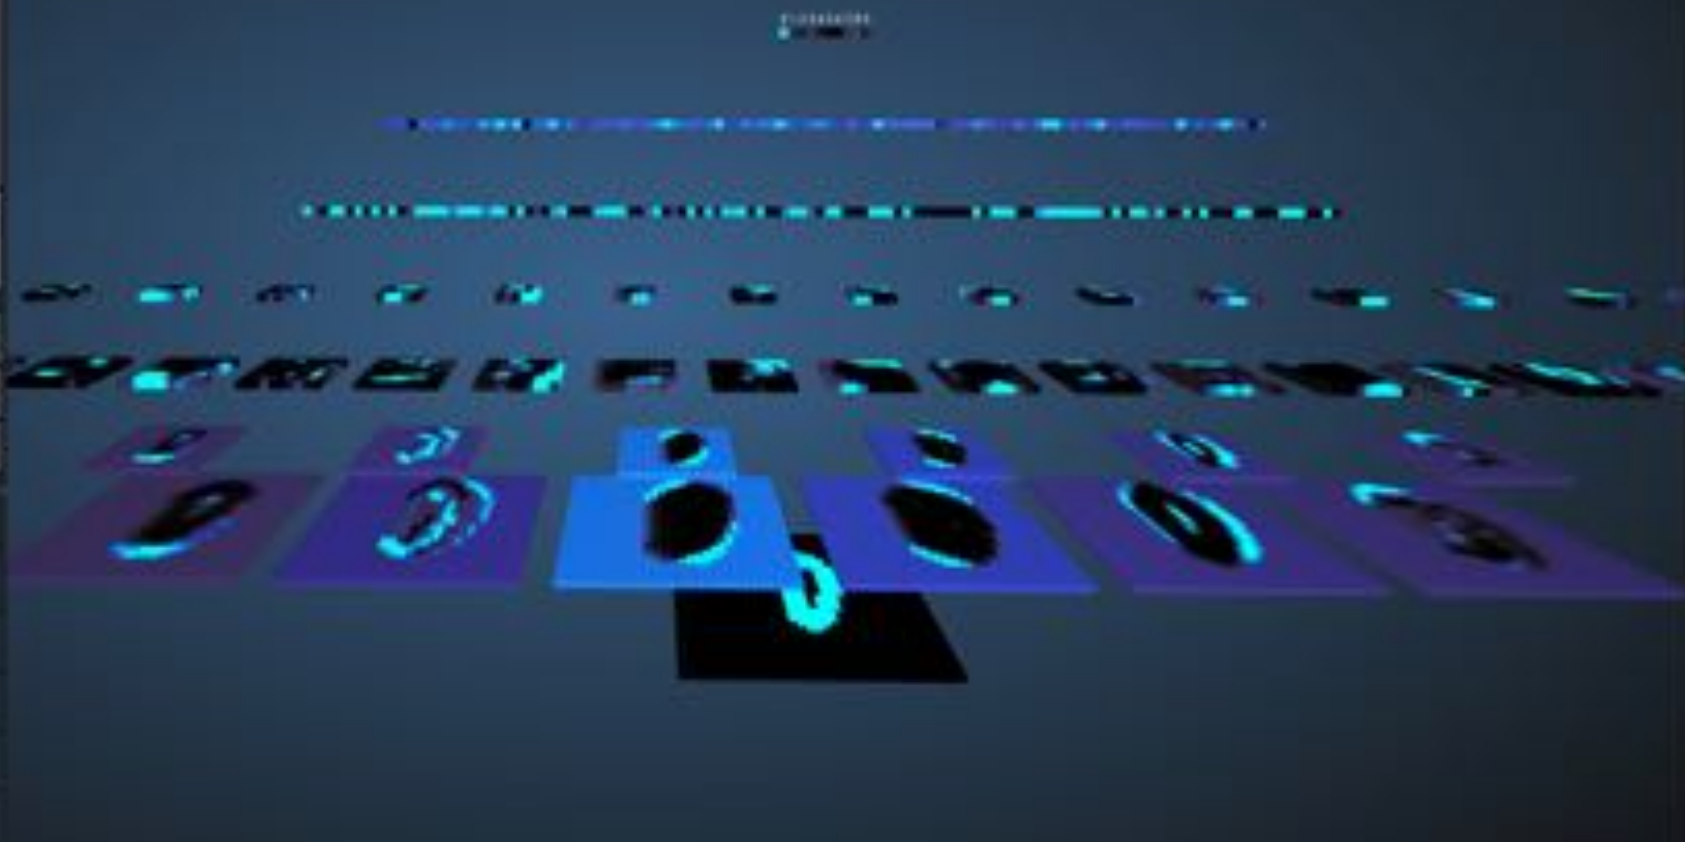
Downsampled drawing: 

First guess:

Second guess:

**Layer visibility**

Input layer	Show
Convolution layer 1	Show
Downsampling layer 1	Show
Convolution layer 2	Show
Downsampling layer 2	Show

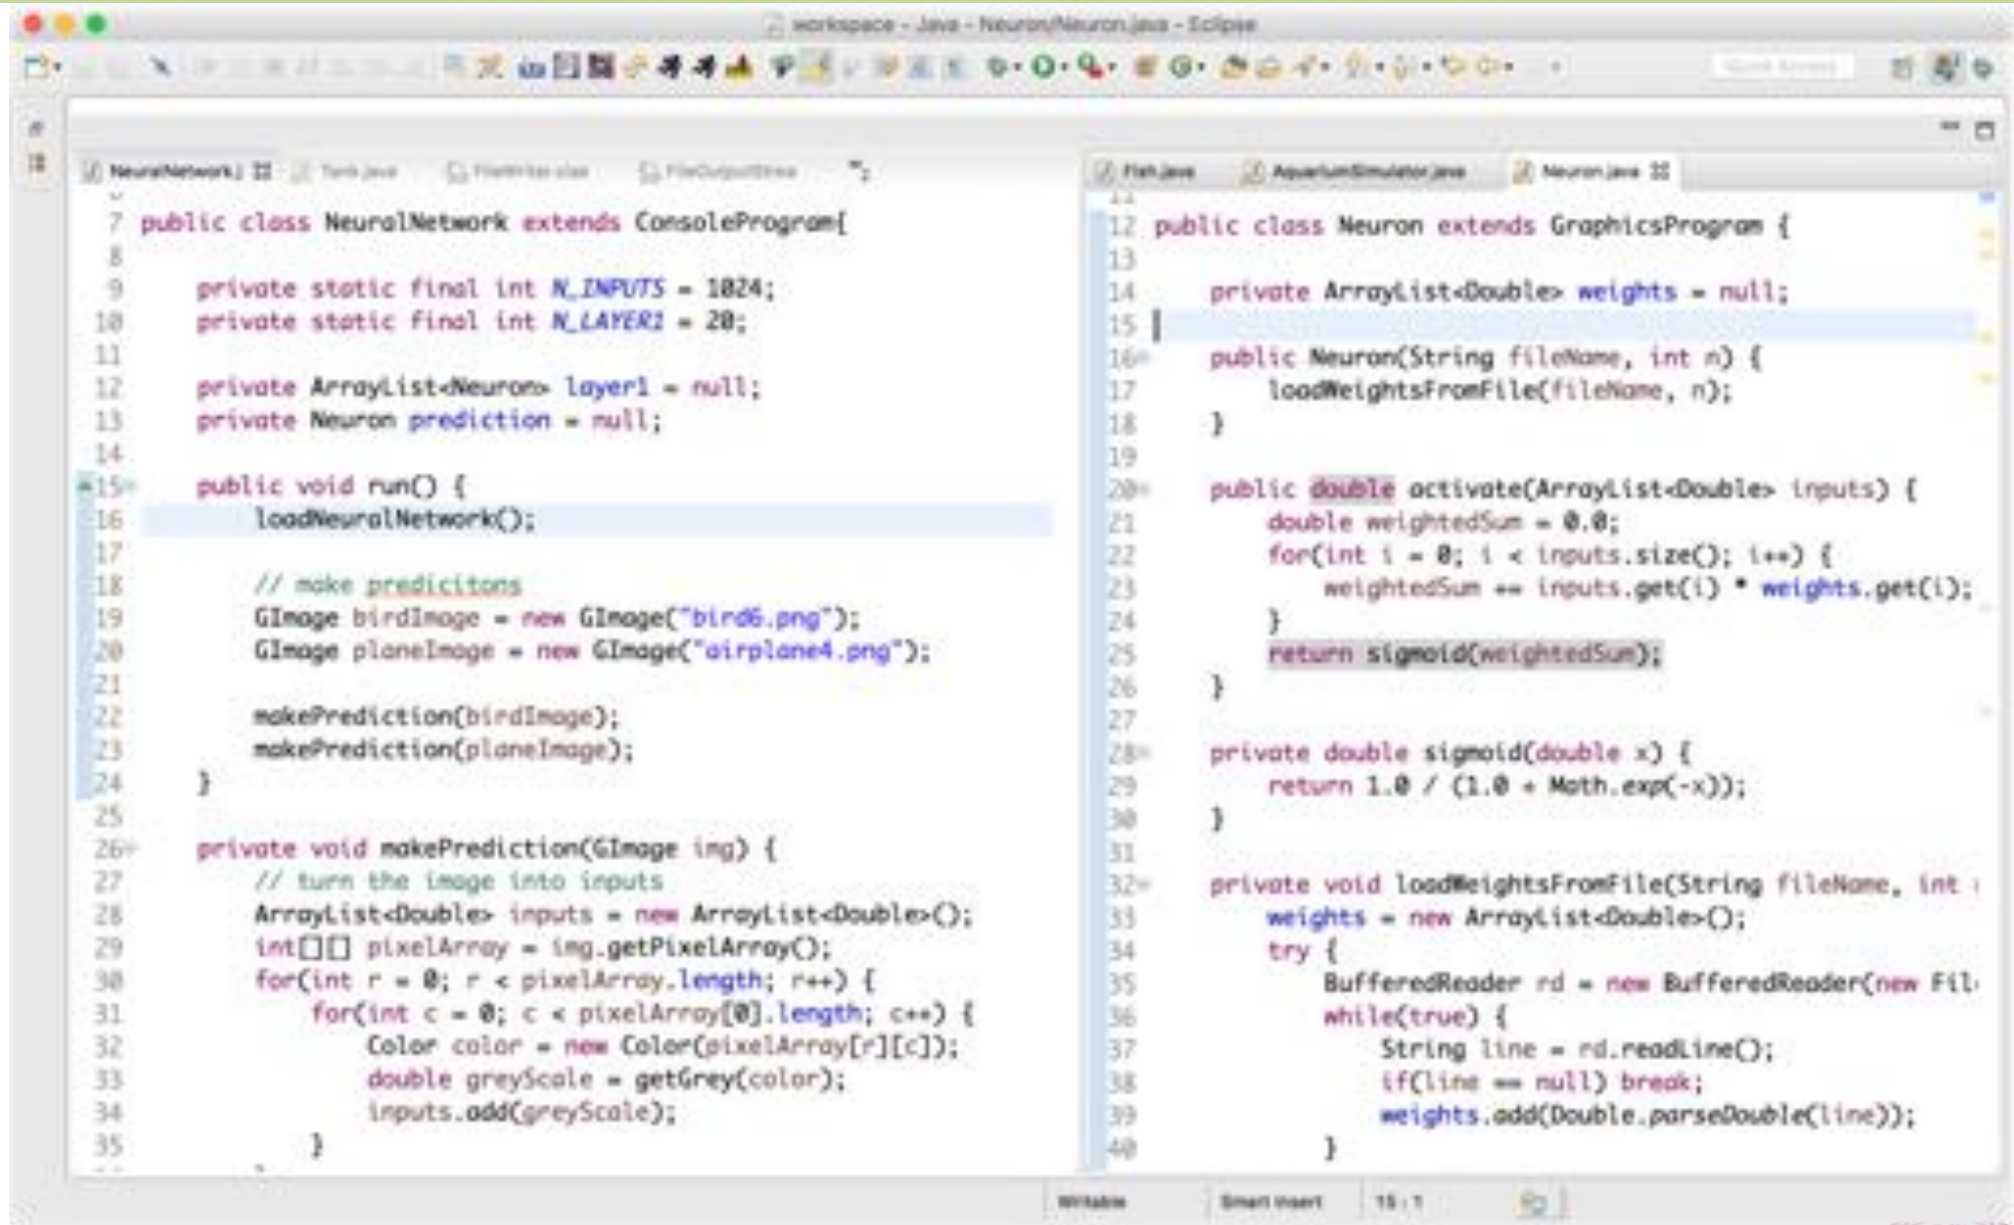


<http://scs.ryerson.ca/~aharley/vis/conv/>

Piech, CS106A, Stanford University



# What Does This Look Like in Code?

The image shows a screenshot of the Eclipse IDE with two Java files open. The left pane shows the code for Neuron.java, and the right pane shows the code for NeuralNetwork.java.

```
Neuron.java
12 public class Neuron extends GraphicsProgram {
13
14     private ArrayList<Double> weights = null;
15
16     public Neuron(String fileName, int n) {
17         loadWeightsFromFile(fileName, n);
18     }
19
20     public double activate(ArrayList<Double> inputs) {
21         double weightedSum = 0.0;
22         for(int i = 0; i < inputs.size(); i++) {
23             weightedSum += inputs.get(i) * weights.get(i);
24         }
25         return sigmoid(weightedSum);
26     }
27
28     private double sigmoid(double x) {
29         return 1.0 / (1.0 + Math.exp(-x));
30     }
31
32     private void loadWeightsFromFile(String fileName, int n
33         weights = new ArrayList<Double>();
34         try {
35             BufferedReader rd = new BufferedReader(new Fil
36             while(true) {
37                 String line = rd.readLine();
38                 if(line == null) break;
39                 weights.add(Double.parseDouble(line));
40         }
}
```

```
NeuralNetwork.java
7 public class NeuralNetwork extends ConsoleProgram{
8
9     private static final int N_INPUTS = 1824;
10    private static final int N_LAYER2 = 20;
11
12    private ArrayList<Neuron> layer1 = null;
13    private Neuron prediction = null;
14
15    public void run() {
16        loadNeuralNetwork();
17
18        // make predictions
19        GImage birdImage = new GImage("bird6.png");
20        GImage planeImage = new GImage("airplane4.png");
21
22        makePrediction(birdImage);
23        makePrediction(planeImage);
24    }
25
26    private void makePrediction(GImage img) {
27        // turn the image into inputs
28        ArrayList<Double> inputs = new ArrayList<Double>();
29        int[][] pixelArray = img.getPixelArray();
30        for(int r = 0; r < pixelArray.length; r++) {
31            for(int c = 0; c < pixelArray[0].length; c++) {
32                Color color = new Color(pixelArray[r][c]);
33                double greyScale = getGrey(color);
34                inputs.add(greyScale);
35            }
}
```



Aside: decomposition

How do we get those weights?

# Chain Rule Down the Network

$$\frac{d \text{ (polyhedron)}}{d \text{ (sphere)}} = \frac{d \text{ (polyhedron)}}{d \text{ (cube)}} \times \frac{d \text{ (cube)}}{d \text{ (sphere)}}$$



# Learning Weights

$$LL(\theta) = y \log \sigma(\theta^T \mathbf{x}) + (1 - y) \log[1 - \sigma(\theta^T \mathbf{x})]$$

---

$$\begin{aligned} \frac{\partial LL(\theta)}{\partial \theta_j} &= \frac{\partial}{\partial \theta_j} y \log \sigma(\theta^T \mathbf{x}) + \frac{\partial}{\partial \theta_j} (1 - y) \log[1 - \sigma(\theta^T \mathbf{x})] \\ &= \left[ \frac{y}{\sigma(\theta^T x)} - \frac{1 - y}{1 - \sigma(\theta^T x)} \right] \frac{\partial}{\partial \theta_j} \sigma(\theta^T x) \\ &= \left[ \frac{y}{\sigma(\theta^T x)} - \frac{1 - y}{1 - \sigma(\theta^T x)} \right] \frac{\partial}{\partial \theta_j} \sigma(\theta^T x) \\ &= \left[ \frac{y - \sigma(\theta^T x)}{\sigma(\theta^T x)[1 - \sigma(\theta^T x)]} \right] \sigma(\theta^T x)[1 - \sigma(\theta^T x)] x_j \\ &= [y - \sigma(\theta^T x)] x_j \end{aligned}$$



Artificial Neurons: One of the greatest decompositions of our lifetimes

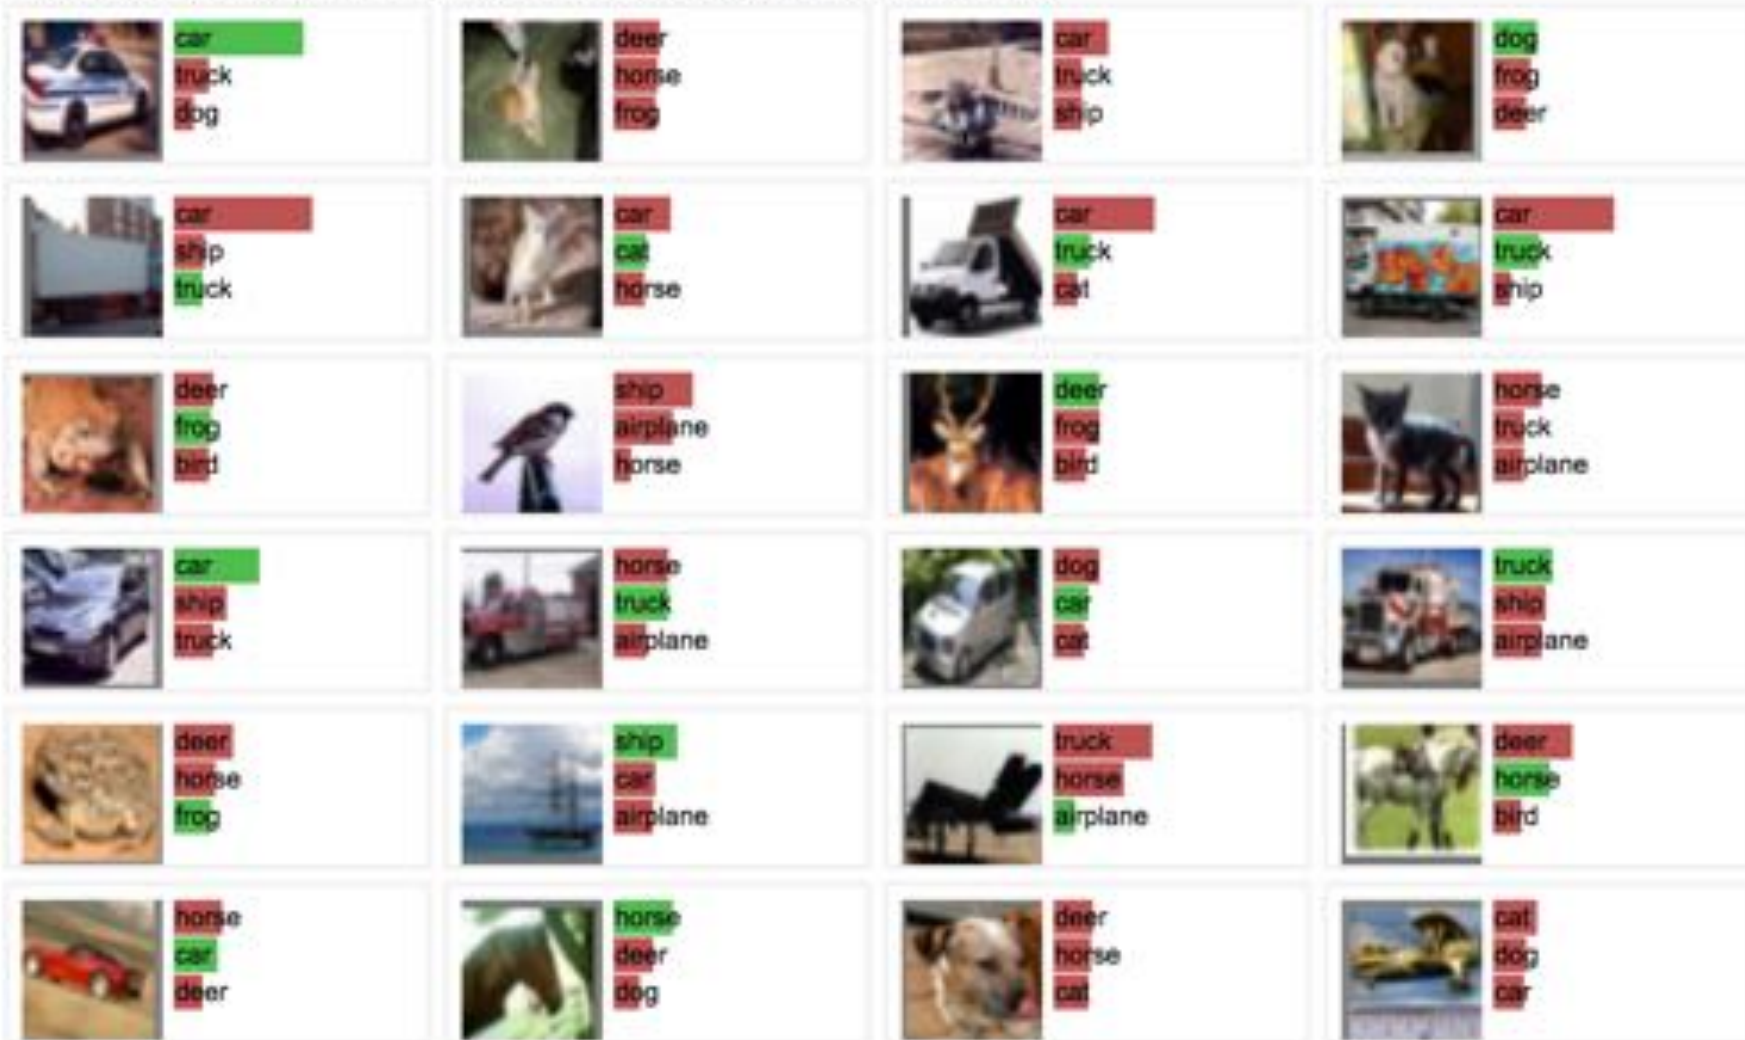
```
model.calculatePartialDerivative(data)
```



```
model.update(data)
```

# Let's Train!

test accuracy based on last 200 test images: 0.2894736842105263



<http://cs.stanford.edu/people/karpathy/convnetjs/demo/classify2d.html>

Piech, CS106A, Stanford University



Like lego pieces

# Visualize the Weights




Training set: Aligned images of faces.




object models



object parts  
(combination  
of edges)




edges



pixels


# GoogLeNet Brain



1 Trillion Artificial Neurons

# GoogLeNet Brain Graph

Multiple,  
Multi class output



22 layers deep





# The Face Neuron



Top stimuli from the test set



Optimal stimulus  
by numerical optimization



# The Cat Neuron



Top stimuli from the test set



Optimal stimulus  
by numerical optimization

**Hire the smartest people in the world**



**Invent cat detector**

# Best Neuron Stimuli

Neuron 1



Neuron 2



Neuron 3



Neuron 4



Neuron 5





# Best Neuron Stimuli

Neuron 6



Neuron 7



Neuron 8



Neuron 9



# ImageNet Classification

22,000 categories

14,000,000 images

Hand-engineered features (SIFT, HOG, LBP),  
Spatial pyramid, SparseCoding/Compression

# 22,000 is a lot!

...

smoothhound, smoothhound shark, *Mustelus mustelus*

American smooth dogfish, *Mustelus canis*

Florida smoothhound, *Mustelus norrisi*

whitetip shark, reef whitetip shark, *Triaenodon obesus*

Atlantic spiny dogfish, *Squalus acanthias*

Pacific spiny dogfish, *Squalus suckleyi*

hammerhead, hammerhead shark

smooth hammerhead, *Sphyrna zygaena*

smalleye hammerhead, *Sphyrna tudes*

shovelhead, bonnethead, bonnet shark, *Sphyrna tiburo*

angel shark, angelfish, *Squatina squatina*, monkfish

electric ray, crampfish, numbfish, torpedo

smalltooth sawfish, *Pristis pectinatus*

guitarfish

rougetail stingray, *Dasyatis centroura*

butterfly ray

eagle ray

spotted eagle ray, spotted ray, *Aetobatus narinari*

cownose ray, cow-nosed ray, *Rhinoptera bonasus*

manta, manta ray, devilfish

Atlantic manta, *Manta birostris*

devil ray, *Mobula hypostoma*

grey skate, gray skate, *Raia batis*

little skate, *Raja erinacea*

...

## Stingray



## Mantaray



0.005%

Random guess

1.5%

Pre Neural Networks

?

GoogLeNet



0.005%

Random guess

1.5%

Pre Neural Networks

43.9%

GoogLeNet

0.005%

Random guess

1.5%

Pre Neural Networks

82.7%

NASNet



# Vision has Social Implications

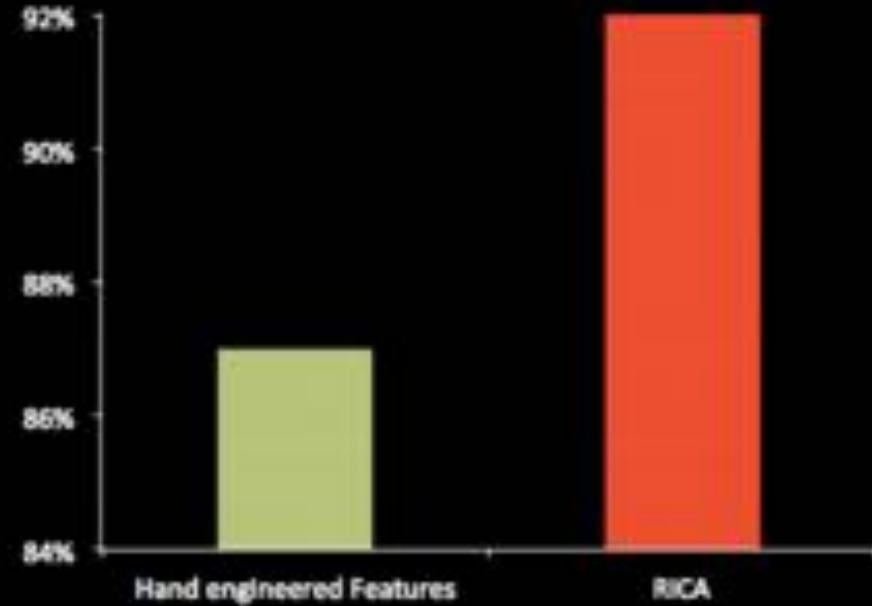
Apoptotic



Viable tumor region




Necrosis



Neural network



# Where is this useful?

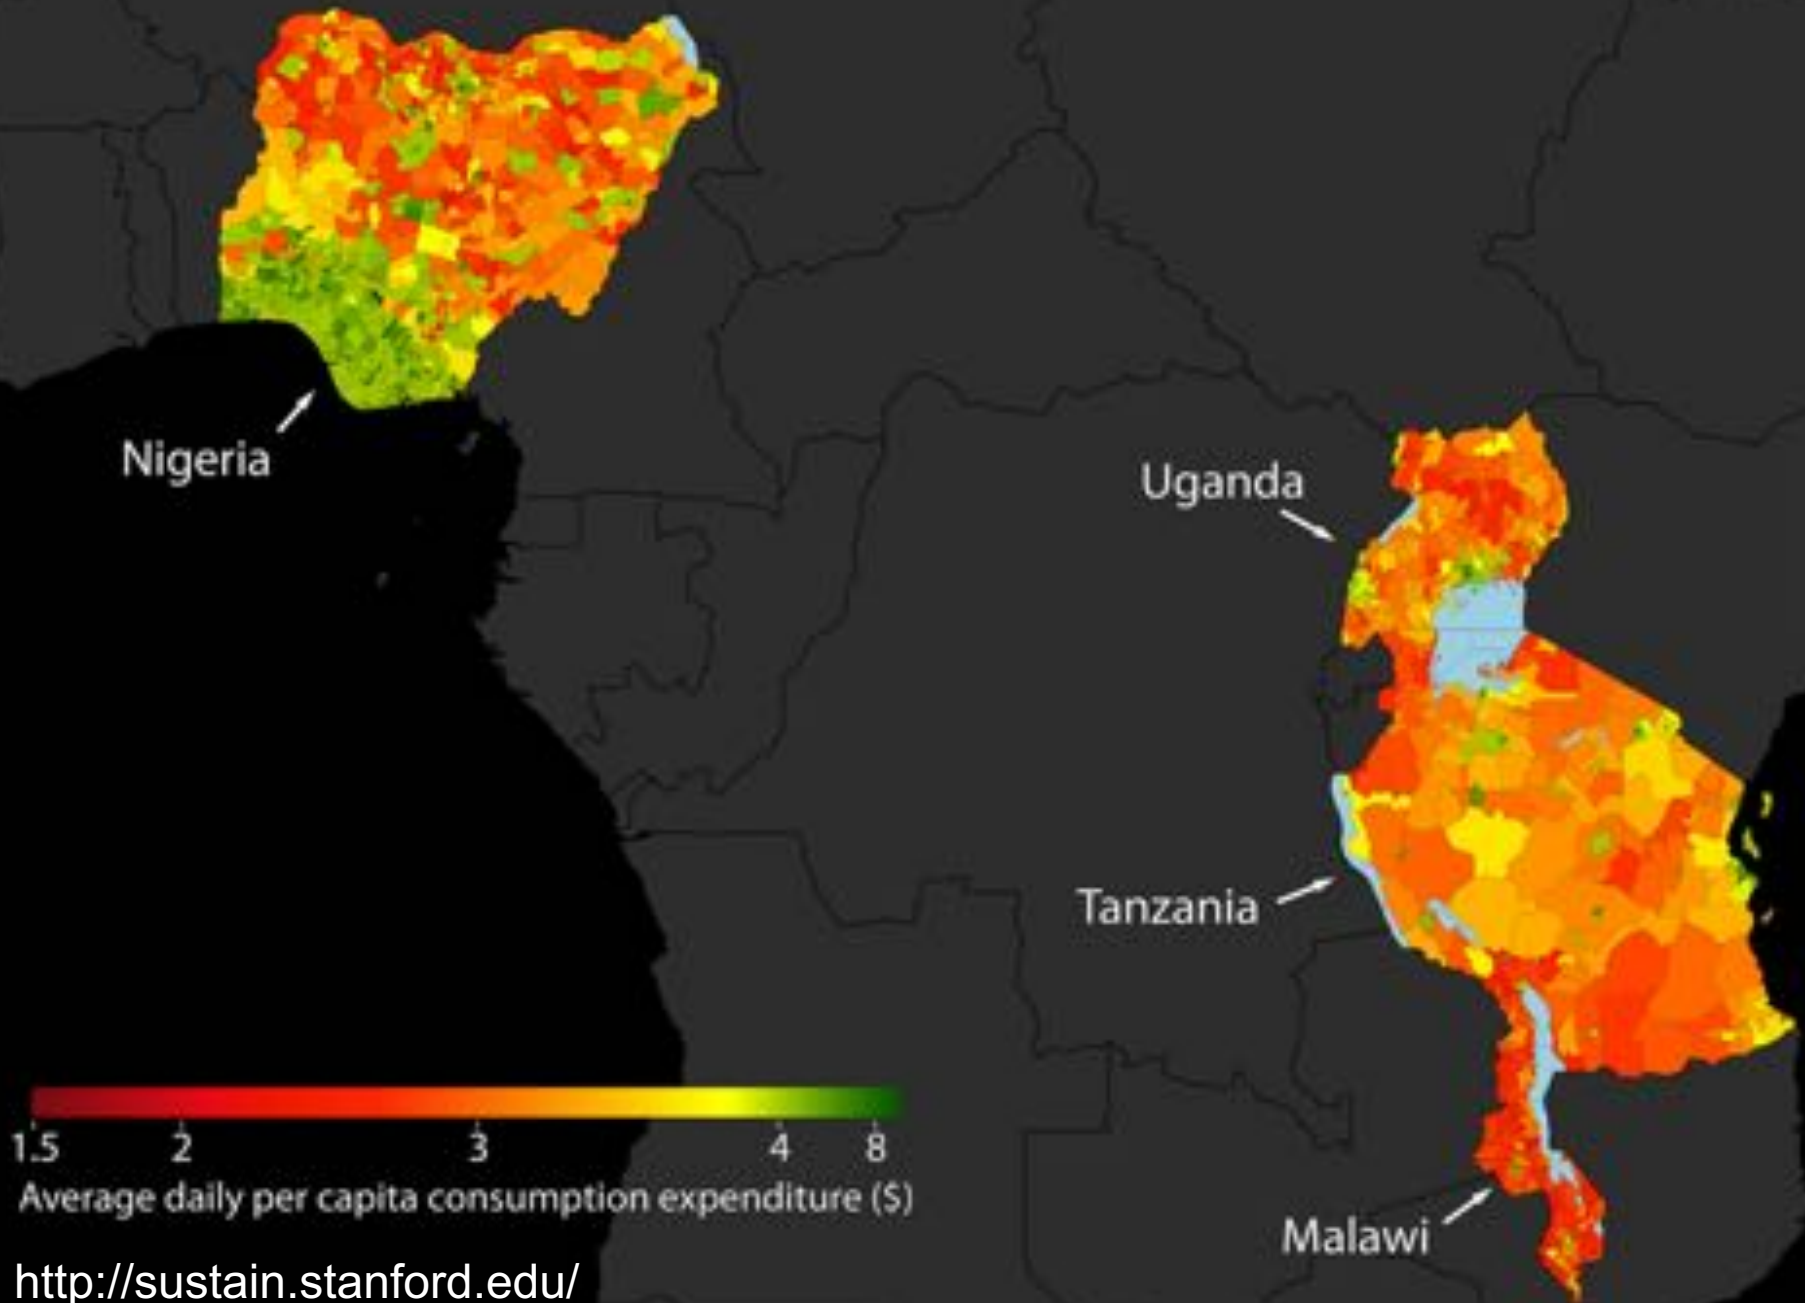


A machine learning algorithm performs **better than** the best dermatologists.

Developed this year, at Stanford.


Esteva, Andre, et al. "Dermatologist-level classification of skin cancer with deep neural networks." *Nature* 542.7639 (2017): 115-118.

# Estimated daily per capita expenditure, 2012-2015



<http://sustain.stanford.edu/>

# Understanding Students



Tl;dr our brain is constantly decomposing




# Told Vision Was 30 Years Out





# Told Speech Was 30 Years Out



Almost perfect...



# What a time to be alive



# Ethics in AI

The end

# Export to JAR

- **JAR: Java Archive.** A compressed binary of a Java program.
  - The typical way to **distribute a Java app as a single file.**
  - Essentially just a ZIP file with Java
- Making a JAR of your project in E
  - File → Export ... →  
Java → **Runnable JAR File**
- *see handout on course web site*

